

Mission Shares Policy on Pleiades

Category: PBS on Pleiades

DRAFT

This article is being reviewed for completeness and technical accuracy.

Mission Directorate shares have been implemented on Pleiades since Feb. 10, 2009. Implementing shares guarantees that each Mission Directorate gets its fair share of resources.

The share to which a job is assigned is based on the GID used by the job. Once all the cores within a Mission Directorate's share have been assigned, other jobs assigned to that share must wait, even if cores are available in a different Mission Directorate's share, with the following exception:

When a Mission Directorate is not using all of its cores, other Mission Directorates can borrow those cores, but only for jobs that will finish within 4 hours. When part of the resource is unavailable, the total number of cores decreases, and each Mission Directorate loses a proportionate number of cores.

You can display the share distribution by adding the "-W shares==" option to the qstat command. For example:

```
%qstat -W shares==
```

Group	Share%	Use%	Share	Exempt	Use	Avail	Borrowed	Ratio	Waiting
Overall	100	0	159748	0	960	158788	0	0.01	960
ARMD	24	18	38109	0	29680	8429	0	0.78	22512
HEOMD	23	21	36521	0	34312	2209	0	0.94	28416
SMD	51	50	80981	0	80968	13	0	1.00	113920
NAS	2	0	3175	0	0	3175	0	0.00	20240

Mission shares are calculated by combining the mission's HECC share of the shared assets combined with the mission-specific assets. The mission shares on Oct 3, 2011 are shown in the second column of the above display. The amount of resources used and borrowed by each mission and resources each mission is waiting for are also displayed.

An in-house utility, *qs*, provide similar information with details that break the resources into the Harpertown, Nehalem-EP and Westmere-EP processor types and is available at </u/scicon/tools/bin/qs>.

The -h option of qs provides instructions on how to use it:

```
% /u/scicon/tools/bin/qs -h
```

```
usage: qs [-u] [-n N] [-b] [-p] [-d] [-r] [-f M,N] [-q N] [-t] [-v] [-h] [--file f]
```

```
-u      : show used resources only; don't show queued jobs
-n N    : show time remaining before N nodes are free
-b      : order segments in bars to help understand borrowing
-p      : plain output: i.e. no colors or highlights
-d      : darker colored resource bars (for a light background)
-r      : use Reverse video for displaying resource bars
-f M,N  : highlight nodes for jobs that finish in <= M minutes
          and <= N minutes [default M=60,N=240]
          (0 turns off highlighting)
-q N    : highlight nodes for jobs queued in last N minutes [3]
          (0 turns off highlighting)
-t      : show time remaining & nodes used for each running job
--file f : reserved for debugging
-v      : (verbose) provide explanation of display elements
-h      : provide this message
```

Here is a sample output file of *qs*:

